

Emergence of slow-switching assemblies in structured neuronal networks

Michael T. Schaub,^{1,*} Yazan N. Billeh,^{2,†} Costas A. Anastassiou,³ Christof Koch,³ and Mauricio Barahona¹

¹*Department of Mathematics, Imperial College London, London SW7 2AZ, U.K.*

²*Computation and Neural Systems Program, California Institute of Technology, Pasadena, California 91125*

³*Allen Institute for Brain Science, Seattle, Washington 98103*

(Dated: July 21, 2015)

Unraveling the interplay between connectivity and spatio-temporal dynamics in neuronal networks is a key step to advance our understanding of neuronal information processing. Here we investigate how particular features of network connectivity underpin the propensity of neural networks to generate slow-switching assembly (SSA) dynamics, i.e., sustained epochs of increased firing within assemblies of neurons which transition slowly between different assemblies throughout the network. We show that the emergence of SSA activity is linked to spectral properties of the asymmetric synaptic weight matrix. In particular, the leading eigenvalues that dictate the slow dynamics exhibit a gap with respect to the bulk of the spectrum, and the associated Schur vectors exhibit a measure of block-localization on groups of neurons, thus resulting in coherent dynamical activity on those groups. Through simple rate models, we gain analytical understanding of the origin and importance of the spectral gap, and use these insights to develop new network topologies with alternative connectivity paradigms which also display SSA activity. Specifically, SSA dynamics involving excitatory and inhibitory neurons can be achieved by modifying the connectivity patterns between both types of neurons. We also show that SSA activity can occur at multiple timescales reflecting a hierarchy in the connectivity, and demonstrate the emergence of SSA in small-world like networks. Our work provides a step towards understanding how network structure (uncovered through advancements in neuroanatomy and connectomics) can impact on spatio-temporal neural activity and constrain the resulting dynamics.

AUTHOR SUMMARY

Neural networks display a wide range of spatio-temporal behaviors. Understanding how this complex orchestration of neuronal firing activity is determined by the structure of the network (*i.e.*, its wiring) is an important step towards comprehending how neural computation is manifested, especially given the growing experimental access to temporal recordings and connectomics. Here we investigate the link between network structure and the dynamics of neuronal assemblies in the context of leaky-integrate-and-fire (LIF) networks. We show how structural features in the wiring of the network can introduce additional time-scales to the dynamics, and how such structured wiring can lead to spatio-temporally segregated, coherent activity of groups of neurons. Using rate models we gain insight into how such spatio-temporal dynamics emerge as a direct consequence of the spectral properties of the network, and use this understanding to examine further circuit topologies that enable phenomenologically similar behavior, yet rely on fundamentally different connectivity and functional patterns.

I. INTRODUCTION

Neuronal ensembles exhibit a broad repertoire of activity patterns. Such dynamics are governed by a time-evolving network of synaptic connections with an intricate, yet structured, organization. Due to the advancement of connectomics, our knowledge about such networks is rapidly growing, and increasingly detailed maps of neuronal wiring are becoming available. In parallel, modern recording techniques, such as calcium imaging and multi-electrode arrays, allow neuroscientists to monitor the activity from thousands of neurons simultaneously, with recordings from entire brains at single neuron resolution becoming technologically feasible [1–3]. The observed dynamics of neural networks exhibit an interplay of structured spatio-temporal scales, which underpin a wide range of cognitive functions [4].

The idea that neuronal group activity induced by network structure is at the core of neural computation dates back at least to the work of Hebb [5], who hypothesized that the transient activity of groups of neurons (so called *cell assemblies*) is the currency of information processing [4, 6]. This notion is supported by recent experiments showing that reciprocal connections between neurons occur above chance level [7, 8], especially if neurons receive common inputs [9, 10]. In the case of the visual system, for instance, excitatory neurons with similar response features tend to be more connected to each other [11, 12]. Moreover, studies have demonstrated that neurons exhibit layer-specific connectivities within rodent sensory cortex [13] and neocortex [14]. In addition, organized architectures have been observed to occur at multiple hier-

* michael.schaub09@imperial.ac.uk; equal contribution

† ybilleh@caltech.edu; equal contribution

archical scales [15–18] and in non-mammalian organisms [19]. These findings suggest that cortical regions contain well-defined subnetworks. However, the underlying question is whether given a network topology, we can predict the potential of the network to sustain structured spatio-temporal activity. Such questions are not only of interest for network dynamics, but have also implications for memory formation and learning, since neural networks undergo topological changes over time due to plasticity [20, 21].

Recently, it has been shown computationally (see e.g. Ref. [22]) that leaky-integrate-and-fire (LIF) networks with equal excitatory and inhibitory connection net strengths (*i.e.* balanced [23]) yet with clustered excitatory connections, can exhibit prolonged heightened group activity, with the activity transitioning between groups in the network (Fig. 1A). Here we characterize the emergence of such slow-switching segregated dynamics in balanced LIF networks as a result of the network connectivity. Specifically, we find that the spectral properties of the synaptic weight matrix (*i.e.*, the existence of an eigenvalue gap and a block-localized dominant subspace) provide a criterion to predict the appearance of such activity in the network. We then use simple linear rate models to gain insight into the mechanisms underpinning the origin of such dynamics in structurally clustered LIF networks. Using these insights, we construct novel LIF topologies that display slow-switching group activity with distinct properties: involving both inhibitory and excitatory neurons; exhibiting multiple slow time-scales; as well as demonstrating the possibility of such dynamics in networks with no obvious clustered connectivity, such as small-worlds. Finally, we discuss briefly possible implications of the different wiring schemes for neural computation.

II. RESULTS

A. Slow-switching assemblies in LIF networks with clustered excitatory connections: spectral insights

Clustered excitatory topologies in a balanced LIF network can lead to dynamics in which localized high activity states transition between assemblies of neurons within the network [22]. This is illustrated in Figure 1A, where the dynamics of an unstructured and a balanced clustered network with 20 groups are shown side by side (see Materials and Methods for a description of the networks). Hereafter, we will refer to such activity as *slow-switching assembly* (SSA) dynamics. Visually, SSA dynamics manifests itself as bands of increased activity in the raster plots, and can be statistically quantified from the resulting spike-train dynamics *a posteriori* (see Eq. (18) in Materials and Methods). Ideally, however, we would like to establish *a priori*, solely from the given connectivity, the possibility of such dynamical patterns emerging.

The full dynamics of LIF networks are notoriously




Figure 1. Network dynamics and eigenvalue spectra of two LIF networks: one with uniform synaptic connections (left), and one with 20 groups of clustered excitatory connections (right). To create the clustered network, excitatory neurons were partitioned into groups (with in-group connection probability p_{in}^{EE} and out-group connection probability $p_{out}^{EE} < p_{in}^{EE}$) while keeping the average connectivity constant (see Materials and Methods and Ref. [22]). The ratio $R_{EE} = p_{in}^{EE}/p_{out}^{EE}$ controls the strength of the excitatory clustering. **A** Visualization of the network topologies (top) and exemplars of raster plots (bottom). The dynamics of the clustered network exhibit the banded structure associated with slow-switching group activity. The magnitude of the activity can be characterized statistically *a posteriori* from the data through the spike-rate variability metric \hat{S} , defined in Materials and Methods Eq. (18), as discussed in the text. In this case, the unclustered network has $\hat{S} = 0.035$ while the clustered network has a much larger value $\hat{S} = 8.23$. **B** Eigenvalue spectra of the network weight matrices W . The weighted connectivity matrix of the clustered network exhibits a clear eigengap $\Delta\lambda$ separating the 19 eigenvalues with largest real parts from the cloud of eigenvalues in the bulk. There is no such eigengap for the unclustered network. As indicated by the two arrows, both matrices have a pair of complex conjugate eigenvalues associated with the (damped) global activation modes of the networks characteristic of balanced networks (see text and Ref. [24]).

difficult to analyze due to their inherent non-linearity; hence an exact analytical treatment of the dynamical evolution for an arbitrary clustered topology is essentially intractable. However, two concepts from spectral graph theory and linear systems provide valuable insights: (i) for symmetric, non-negative connectivity ma-

trices, it can be shown that a modular network structure implies a gap in the spectrum of the graph (i.e., in the set of eigenvalues of the weight matrix) [25, 26], as well as the block-localization of the associated eigenvectors on the modules of the network (noting that isolated eigenvalues may also be the result of other features [27]); (ii) for linear systems, a gap in the spectrum of the graph results in a *separation of time scales* in the dynamical process [24, 28, 29]. This relation between the modular structure, the eigenvalues and associated eigenvectors, and linear network dynamics can be used to discover modular structures in graphs from a dynamical perspective [30–33].

In fact, the weighted connectivity matrices of unclustered and clustered LIF networks [22] display different spectral characteristics. Figure 1B shows the spectra of two networks with 1600 excitatory and 400 inhibitory neurons with unclustered (left) and clustered (right) topologies. In the unclustered case, we find the expected circular distribution of eigenvalues, which follows from the properties of random graphs [34, 35], although the presence of groups of neurons with different cardinalities and variances means that the eigenvalue distribution is not completely uniform on the circle [34]. Note also that the balanced construction of the LIF network, with a marginally larger inhibitory input for each neuron in order to keep the network stable, leads to the existence of one pair of complex conjugate eigenvalues which lies separate from the main bulk of the spectrum (black arrows in Figure 1B). This eigenvalue pair is associated with the global activation mode of the network (as explained below in Section II B and in Ref. [24]). As shown in Figure 1A, this unclustered network exhibits the expected asynchronous, unstructured neuronal spiking dynamics.

In contrast, the LIF network with clustered excitatory neurons displays banded SSA dynamics. Spectrally, its weight matrix exhibits a clear gap $\Delta\lambda$ along the real axis of its spectrum (Figure 1B, right) and, as shown below in Section II B 3 (Fig. 5), the associated Schur vectors also exhibit a measure of structural block-localization. We remark that, in general, this should not be expected *a priori*, since the weight matrices are *asymmetric* and include both positive (excitatory) and negative (inhibitory) couplings.

To ascertain the dynamical relevance of the spectral gap, we examined the relation between the clustering strength in the network, defined as the ratio of probabilities of connections inside and outside the neural assemblies ($R_{EE} = p_{in}^{EE}/p_{out}^{EE}$), the spectral gap ($\Delta\lambda$), and the magnitude of the numerically observed SSA dynamics. To quantify the assembly spike-rate variability, we have defined two complementary metrics. First, the metric \hat{S} measures the heterogeneity in the firing rates of the putative cell-assemblies averaged over the simulation (see Eq. (18) in Materials and Methods). A large value of \hat{S} indicates that the average firing rates of the assemblies are diverse, whereas a low \hat{S} indicates that all groups have very similar firing rates at all times and no group




Figure 2. The relationship of observed SSA dynamics with the structural connectivity clustering of the LIF network, R_{EE} , and the spectral gap, $\Delta\lambda$. The presence of SSA dynamics is quantified through the spike-rate variability metric (\hat{S}), which measures the variance of the firing rates of the assemblies normalized by a randomly shuffled bootstrap: \hat{S} increases with increasing SSA activity and $\hat{S} \approx 0$ for completely asynchronous activity (as in the unclustered case in Figure 1A) (see Materials and Methods). **A** Spike-rate variability is plotted as a function of R_{EE} (left) and $\Delta\lambda$ (right) for different network sizes (dots: raw data from simulations, line: mean, shading: standard deviation). Above a certain clustering threshold, SSA emerges and increases as R_{EE} grows; the intensity of the SSA dynamics is in line with the presence of an eigenvalue gap $\Delta\lambda$ in the weight matrix. **B** Spike-rate variability as a function of R_{EE} (left) and $\Delta\lambda$ (right) for a network of 2000 neurons with different numbers of clusters, yielding qualitatively similar results (dots: raw data from simulations, line: mean, shading: standard deviation). **C** Relationship between the clustering strength R_{EE} and the spectral gap $\Delta\lambda$. Observe that R_{EE} is not sufficient to determine $\Delta\lambda$, i.e. $\Delta\lambda$ is influenced by other aspects such as the network size and number of groups.

shows elevated firing. Under SSA dynamics, the variability of firing rates across groups increases in time as the assemblies transition between high and low firing rates. As discussed below in Section II B 4, it is possible that the heightened firing activity is localized in a particular assembly and does not switch between assemblies. This

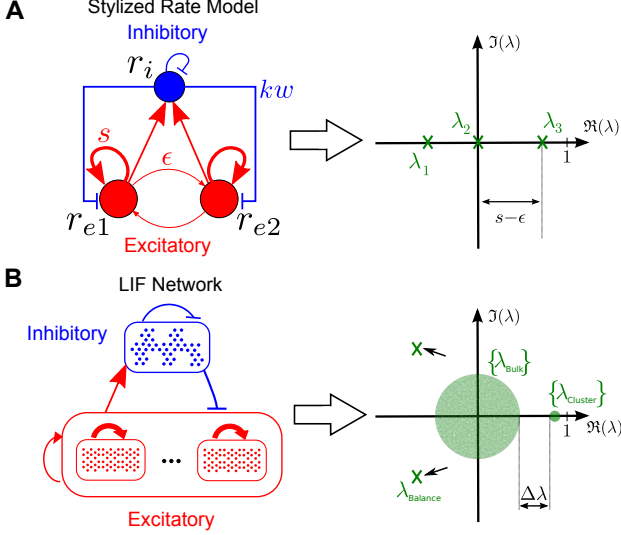


Figure 3. Structure of the eigenspectra of stylized and realistic networks that can exhibit SSA activity. **A** Schematic of the stylized model of a clustered network (1)–(2) and illustration of the corresponding eigenvalues (6). **B** Schematic of a more realistic clustered LIF network with the same wiring scheme on many groups and multiple neurons per group with an illustration of the associated spectrum.

non-switching dynamics where only a group of neurons exhibits elevated firing over the entire simulation time can be detected using the second variability measure \hat{S}_T , defined in Eq. (20). For SSA dynamics, both \hat{S} and \hat{S}_T give similar results (see Fig. S1). In the rest of the paper, we focus on SSA dynamics and mainly use \hat{S} , except when discussing the transition to non-switching behavior in Section II B 4. Definitions of these quantities are given in Materials and Methods.

Figure 2 shows that SSA becomes observable above a threshold of the clustering strength R_{EE} , although the dependence of this threshold on the size of the network and number of clusters does not seem to follow an obvious pattern. On the other hand, our numerics show that the spectral gap provides a more direct indicator of the presence of SSA dynamics in the network. The relationship between R_{EE} and $\Delta\lambda$ (Figure 2C) shows that knowing R_{EE} is not sufficient to determine $\Delta\lambda$ as the spectral gap depends on other factors such as the network size and the number of groups in the network. Together with the examination of the structure of the associated orthogonal subspace (see Section II B 3), such spectral characterization may be used to establish the potential of networks to sustain SSA activity, even if there is no obviously clustered network model known *a priori*.

The observed spectral properties of clustered LIF networks lead us to consider a stylized linear rate model for neuronal activity in the next section, as a means to gain insights into the defining factors of the wiring diagram leading to SSA.

B. Linear rate models and slow localized activity in clustered LIF networks

1. A stylized linear rate model for networks with clustered excitatory neurons

Motivated by the above findings, we proceeded to investigate how much of the observed LIF dynamics can be described by simple rate models. From above, we expect that the spectral properties of the network lead to SSA dynamics that is coherent over, and transitioning between, clusters. Hence we consider simple rate models on coarse-grained networks where each node describes the behavior of a cluster. As will become clear below, this reduced description allows us to identify the mechanism underpinning the onset of SSA activity without having to consider a plethora of statistical and technical detail. Furthermore, our main conclusions can be easily extended to multiple groups of clustered excitatory neurons, or to scenarios with probabilistic connectivity between the nodes (see e.g. Ref.[24]).

To explore this idea in the simplest setting, consider a linear firing rate model with three groups: two groups of clustered excitatory neurons and a group of inhibitory neurons with uniform coupling. This wiring scenario is abstracted in the form of a three-node network (Figure 3A), where each node represents a group of neurons, and a simple rate model is given by the following equations:

$$\tau \frac{d\mathbf{r}}{dt} = -\mathbf{r} + W\mathbf{r} + \boldsymbol{\xi} = -(I - W)\mathbf{r} + \boldsymbol{\xi}. \quad (1)$$

Here $\mathbf{r} = (r_{e1}, r_{e2}, r_i)^T$ are the firing rates of the neuron groups; W is the synaptic connectivity weight matrix; and $\boldsymbol{\xi}$ is a random input to the network. The firing rates are measured relative to some baseline activity, and may be positive or negative.

To describe this topology with two excitatory clusters we define a connectivity matrix W of the form:

$$W = \begin{pmatrix} s & \epsilon & -kw \\ \epsilon & s & -kw \\ w/2 & w/2 & -kw \end{pmatrix}, \quad (2)$$

$$s > \epsilon > 0, \quad k \geq 1, \quad w = (s + \epsilon). \quad (3)$$

The first inequality in (3) guarantees that the *coupling strength* is positive: $s - \epsilon > 0$, i.e., the cross-coupling between the excitatory groups is weaker than the intra-group coupling, as would be expected from the notion of a cluster. The second inequality in (3) ensures a balanced network (while keeping $k \approx 1$), such that each group has at least the same amount of inhibitory and excitatory input: $\sum_j W_{ij} = (1 - k)w \leq 0, \forall i$.

To understand the dynamical behavior of this system, we evaluate the spectral properties of the connectivity matrix W [24, 29, 36]. Note that we are dealing with a non-normal system, in which the eigenvectors may not be orthogonal and thus will provide a possibly misleading description of the dynamics [37]. We therefore consider

the Schur decomposition of W , which yields an orthonormal basis of the system [24, 36], i.e., it finds an orthogonal matrix U and an upper triangular matrix Q such that $W = UQU^T$. Here U contains an orthonormal basis while Q contains the eigenvalues of W on the diagonal and non-zero elements above the diagonal accounting for effective feedforward connections between the different modes [24, 36].

For the matrix (2), the associated upper triangular Schur form Q is:

$$Q = \begin{pmatrix} -w^+ & w_{ff} & 0 \\ & 0 & 0 \\ & & s - \epsilon \end{pmatrix}, \quad (4)$$

$$w^+ = w(k-1), \quad w_{ff} = (k+1/2)(w+\epsilon),$$

and the matrix containing the orthonormal (Schur) basis is $U = (\mathbf{u}_1 \mathbf{u}_2 \mathbf{u}_3)$ given by

$$\mathbf{u}_1 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \quad \mathbf{u}_2 = \sqrt{\frac{2}{3}} \begin{pmatrix} 0.5 \\ 0.5 \\ -1 \end{pmatrix}, \quad \mathbf{u}_3 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}. \quad (5)$$

Each of the Schur vectors \mathbf{u}_i represents a pattern of rates, a particular mode of firing on the network. The first mode \mathbf{u}_1 is associated with an overall mean firing pattern, while the second mode \mathbf{u}_2 corresponds to a relative difference in firing between inhibitory and excitatory groups. Note that in both cases the firing patterns of the two groups of excitatory neurons are aligned. In contrast, the third mode \mathbf{u}_3 describes an antagonistic activity localized on the excitatory neuron groups (i.e., when the firing rate of one excitatory group increases the other decreases, and vice versa) with the inhibitory node remaining unaffected at a baseline firing rate, precisely in line with the SSA behavior observed in the full LIF network (Fig. 1A). Observe also that, while the first two modes are coupled via an off-diagonal term in Q (i.e., an effective feedforward connection [24, 36]), the switching behavior described by the third mode is uncoupled from the other system modes.

The structure of the Schur decomposition provides us with insight into the dynamics of the model (1). From (4), the eigenvalues of W are:

$$\lambda_1 = -w^+ \leq \lambda_2 = 0 < \lambda_3 = s - \epsilon, \quad (6)$$

whence the eigenvalues governing (1) are $\eta_1 = -1 - w^+$, $\eta_2 = -1$ and $\eta_3 = -1 + (s - \epsilon)$. For the system to be linearly stable, we require $\eta_3 < 0$, which implies that the clustering strength in a stable system is constrained to be

$$0 < s - \epsilon < 1. \quad (7)$$

Without an external input, the three modes decay exponentially with time constants $\tau_i = -1/\eta_i = 1/(1 - \lambda_i)$. Therefore, it is possible to slow down the time scale associated with the SSA mode \mathbf{u}_3 by increasing the clustering strength $s - \epsilon \rightarrow 1$. The perturbations introduced by




Figure 4. SSA dynamics can be achieved by changing solely the synaptic strengths even when the topological connections are kept uniform. To create clustered networks, the ratio $W_{EE} = w_{in}^{EE}/w_{out}^{EE}$ was varied, where w_{in}^{EE} and w_{out}^{EE} refer to the in-group and out-group synaptic weights, respectively. The spike-rate variability \hat{S} measuring the intensity of SSA dynamics is shown as a function of W_{EE} (left - dots: raw data from simulations, line: mean, shading: standard deviation) and the spectral gap $\Delta\lambda$ (right). Note that the connection probability is uniform for all the simulations, i.e. $R_{EE} = 1$, and only the clustering of the weights is varied.

the random input induce mode-mixing and, in particular, break the symmetry between the firing rates of the two groups of excitatory neurons. These perturbations excite the SSA mode \mathbf{u}_3 , which will decay with time scale $1/(1 - \lambda_3)$ towards the solution with uniform rates on all groups ($\mathbf{r}^* = 0$).

In conclusion, the stronger the clustering (i.e., the larger $s - \epsilon < 1$), the slower the decrease of the asymmetric transients, and the more prominent SSA dynamics becomes. Note that this slow time scale in the activity patterns is directly associated with the eigengap $\lambda_3 - \lambda_2 = s - \epsilon$ (Fig. 3), which is dictated by the clustering strength, in agreement with our observations relating the presence of SSA with $\Delta\lambda$ in Figure 2. The insights gained from this simple model have testable implications for LIF networks and lead to further ideas for neuro-physiological network wiring structures, which we explore in the following sections.

2. The linear rate model and the full dynamics of clustered LIF networks

Our rate model can be extended to networks with c groups and the overall structure of the eigenspectrum of W will essentially retain its features (Figure 3; see [24] for a related discussion). First, there will be a ‘negative’ eigenvalue (λ_1 in (6) related to the pair λ_{Balance} in Fig. 3B), which reflects the fact that the network is balanced and stable, i.e. the associated global activity mode decays. Second, there will be a bulk of ‘small’ eigenvalues centered around the origin (λ_2 in (6) and in general the set $\{\lambda_{\text{Bulk}}\}$ in Fig. 3B), which stem from the random connectivity present in the network—a random matrix has a circular eigenvalue distribution around the origin). Finally, there will be an eigenvalue gap separating a small

set of $c - 1$ ‘large positive’ eigenvalues (λ_3 in (6) and $\{\lambda_{\text{Cluster}}\}$ in Fig. 3B) from the bulk of the spectrum. These are the ‘slow’ eigenvalues associated with an orthogonal Schur subspace with a block structure coherent with the clusters, and are thus responsible for the appearance of a new time scale in the system originating the observed SSA dynamics. This structure of the eigenvalues can be seen in the spectrum of the LIF network in Figure 1B.

Inspired by the analysis of the linear rate model, we investigate how the results translate to the full nonlinear LIF dynamics. First, it is clear from above that only the difference between the effective intra- and inter-group coupling strengths is essential for the observed slow activity. Such difference can be the result of clustered topological connectivity as in Ref. [22], but can be equally achieved keeping the topology uniformly connected and tuning the synaptic weights. To assess this possibility, we conducted numerical simulations of LIF networks with uniform connection probabilities between excitatory neurons, yet with larger intra-group synaptic weights while keeping the average weight constant (see Materials and Methods). Figure 4 shows that tuning the weights towards a clustered structure (as given by the weight ratio $W_{EE} = w_{\text{in}}^{EE}/w_{\text{out}}^{EE}$) also leads to the emergence of SSA dynamics linked to a spectral gap. Although unsurprising from a linear systems perspective, this result confirms the applicability of the spectral characterization to nonlinear LIF networks and, specifically, establishes its relevance for SSA dynamics in networks where only synaptic weights can be tuned.

3. The block-localization of the dominant linear subspace

We now consider a critical factor in the emergence of SSA, namely, that the leading Schur vectors of the weight matrix of LIF networks exhibit a measure of consistent block-localization on the groups of neurons which then exhibit cell assembly dynamics (Figure 5). Due to the non-normality outlined above [24, 36], we consider orthogonal Schur vectors and not eigenvectors. A typical Schur vector corresponding to one of the eigenvalues in the leading group $\{\lambda_{\text{Cluster}}\}$ exhibits a coherent pattern on each of the groups of the network, such that all neurons within a particular group are driven uniformly. This block-uniformity results in the grouped dynamics that characterizes SSA. Representative patterns are shown in Figure 5A. In contrast, the Schur vectors from the bulk of the spectrum $\{\lambda_{\text{Bulk}}\}$ show no specific pattern of localization on any group of neurons. Note that the Schur vectors show no localization pattern on the inhibitory neurons either, in line with our analysis. The block-localization of the leading Schur vectors constitutes an additional spectral characterization to assess the emergence of SSA dynamics from the weight matrix alone.

In Figure 5, we quantify how the patterns observed in the SSA dynamics align with the dominant Schur vectors

of the weight matrix. To do this, we first performed a principal component analysis (PCA) of the simulated firing rates of the full LIF network and extracted the top activation patterns $\{\mathbf{p}_1, \dots, \mathbf{p}_c\}$ corresponding to the first principal components (see Figure 5B and Materials and Methods). To measure the alignment of the data patterns $\{\mathbf{p}_i\}$ with the dominant Schur vectors $\{\mathbf{u}_1, \dots, \mathbf{u}_c\}$ of the weight matrix W corresponding to the leading eigenvalues above the gap, we computed the *first principal angle* θ between these two subspaces: the smaller the angle, the closer the alignment [38, 39]. Figure 5D shows that the dominant subspace of the Schur vectors of the weight matrix (W) are indeed aligned with the observed patterns and consistent with the embedded groups of neurons in the network. Finally, Figure 5E shows again that tuning the synaptic weights to cluster the network while keeping the topology uniformly connected renders a similar outcome.

4. Increasing the clustering beyond the linearly stable regime: one dominant assembly

Until now, our linear analysis has concentrated on the relevant regime for SSA, where the system is linearly stable (7): $\lambda_3 = s - \epsilon < 1$. However, it is possible to increase the clustering strength ($s - \epsilon$), so that the intra-cluster connections are much stronger than the inter-cluster connections, and the linear system has an unstable eigenvalue. For the LIF network, if we increase R_{EE} or W_{EE} so that $\lambda_{\text{max}} \in \{\lambda_{\text{Cluster}}\} \geq 1$, then the associated localized firing mode, once excited, will activate itself. This regime can thus lead to a ‘winner takes all’ solution, where one cell assembly dominates the firing of the whole network and suppresses all other neurons [40].

From the perspective of the observed dynamics of the LIF network, an increase of the clustering strength (R_{EE} or W_{EE}) initially leads to the development of the eigen-gap for $\{\lambda_{\text{Cluster}}\}$ and, the emergence of SSA dynamics. However, when the clustering strength becomes too large (and the largest eigenvalue goes beyond 1), a single assembly starts to dominate the firing pattern and the firing variability is reduced across time, an effect which is captured by a strong reduction in \hat{S}_T , as illustrated in Figure 6. This behavior, which has been observed previously [22], is thus also closely linked to the spectral properties of the system. Note that the linear condition $\lambda_{\text{max}} \geq 1$ is only indicative: the non-linearity and boundedness of LIF dynamics can lead to (non-linearly stable) SSA dynamics beyond this condition.

C. Beyond clustered excitatory neurons: SSA dynamics involving both excitatory and inhibitory neurons.

So far we showed that a linear rate model can provide valuable insights into the mechanisms underpinning SSA

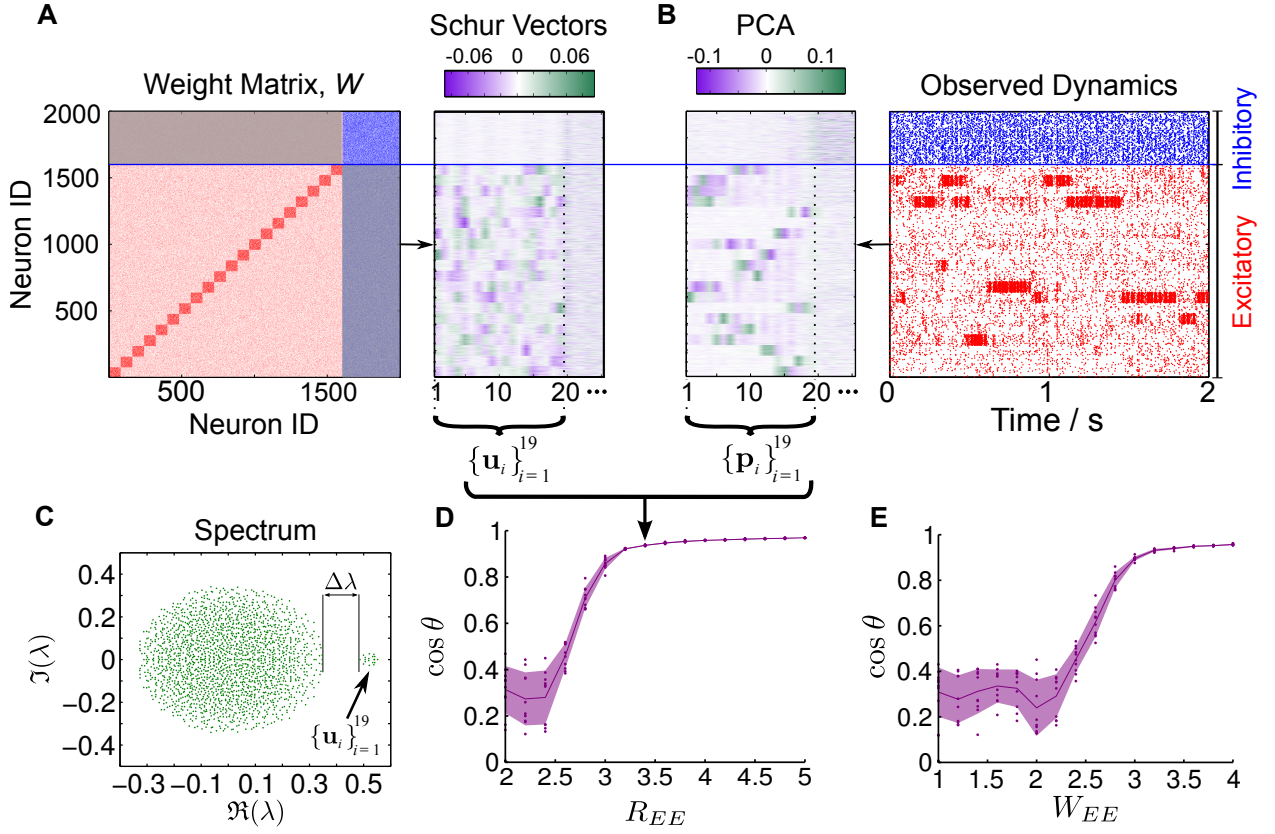


Figure 5. Unraveling the association between the leading Schur vectors of the weight matrix and the observed SSA activity. **A** Illustration of a weight matrix of the $N = 2000$ neuron LIF network with $c = 20$ groups of excitatory neurons ($R_{EE} = 3.4$) shown together with the heatmap of the real parts of its first 25 Schur vectors. The first 19 Schur vectors exhibit block-uniform patterns relatively constant within each of the 20 groups of neurons, whereas no such pattern is observed for the other Schur vectors. Note that there is no pattern discernible over the inhibitory neurons. The leading Schur vectors correspond to the cloud of ‘slow’ eigenvalues above the gap, as indicated in **C**, and span a patterned dominant subspace that induces grouped dynamics in the network. **B** A long simulation (80s) of the LIF network dynamics (only the first 2s shown) was analyzed using PCA and the first 25 principal components (PCs) are shown. Reflecting the banded structure of the simulated dynamics, the leading PCs also show a block-patterned structure consistent with the neuronal groups. **C** On the spectrum of W , we indicate the group of leading eigenvalues above the gap associated with the dominant subspace. **D** The alignment between the dominant Schur subspace of the W matrix and the subspace of the strongest principal components is measured by the first principal angle θ (21). Above a threshold of the clustering strength R_{EE} , both subspaces become highly aligned in line with the observations in Fig. 2 (dots: raw data from simulations; line: mean; shading: standard deviation). **E** The same effect is observed when the clustering is introduced in the weights by varying W_{EE} , as in Fig. 4.

dynamics in networks with clustered excitatory neurons. The crucial point for the emergence of SSA is the separation of time scales, dictated by the splitting of the leading eigenvalues of the weight matrix, together with the block-localization of the associated dominant subspace. These spectral properties can be introduced into LIF networks by entirely different synaptic couplings, and we now consider a mechanism in which the inhibitory neurons are involved more centrally in generating SSA dynamics.

Consider the wiring diagrams presented in Figure 7A-B, representing the simple rate model and its corresponding LIF schematic. In this case, the wiring does not involve clustered coupling between groups of excitatory neurons, but rather relies on preferential connectivity patterns *between* inhibitory and excitatory neuron groups

only. Each group of excitatory neurons activates preferentially an associated group of inhibitory neurons and, in turn, this group of inhibitory neurons feeds back more weakly to its associated group of excitatory neurons (see Materials and Methods for a full description of the network). Therefore the effective functional circuitry consists of both excitatory and inhibitory neurons embedded in a feedback loop and, as a consequence, the inhibitory neurons play an integral role in generating the spatio-temporal dynamics and display SSA dynamics too, as we show below.

This wiring mechanism was suggested by the stylized firing rate model with two coupled pairs of inhibitory/excitatory feedback loops in Figure 7A. This system is described by Eq. (1) with a vector of firing




Figure 6. Effect of large clustering strength on SSA activity. **A** Large values of the clustering R_{EE} lead to linear instability of the SSA dynamics and localization of the activity on one assembly. As measured by the spike rate variability across time (\hat{S}_T), the increase of R_{EE} leads to SSA (signalled by the increased value of \hat{S}_T). If the clustering increases further, \hat{S}_T decreases, as the dynamics becomes dominated by one assembly only. The dots denote raw data from simulations (line: mean; shading: standard deviation). Inset: examples of raster plots for three data points in the three regimes. The analysis corresponds to a clustered LIF network of 1000 neurons. **B** Plot of the eigenvalue with the largest real component λ_{\max} against the clustering strength R_{EE} . The linear condition $\lambda_{\max} > 1$ is a good indicator of the dynamics becoming dominated by one cell assembly.

rates for the four groups $\mathbf{r} = (r_{e1}, r_{e2}, r_{i1}, r_{i2})^T$ and a synaptic weight matrix defined as:

$$W = \begin{pmatrix} w & w & -k\epsilon & -ks \\ w & w & -ks & -k\epsilon \\ s & \epsilon & -kw & -kw \\ \epsilon & s & -kw & -kw \end{pmatrix}, \quad (8)$$

$$s > \epsilon, \quad k \geq 1, \quad w = (s + \epsilon)/2. \quad (9)$$

Similarly to (2), the system is balanced and $s - \epsilon$ captures the clustering strength within the excitatory-inhibitory pairs. The Schur decomposition of W leads to the fol-

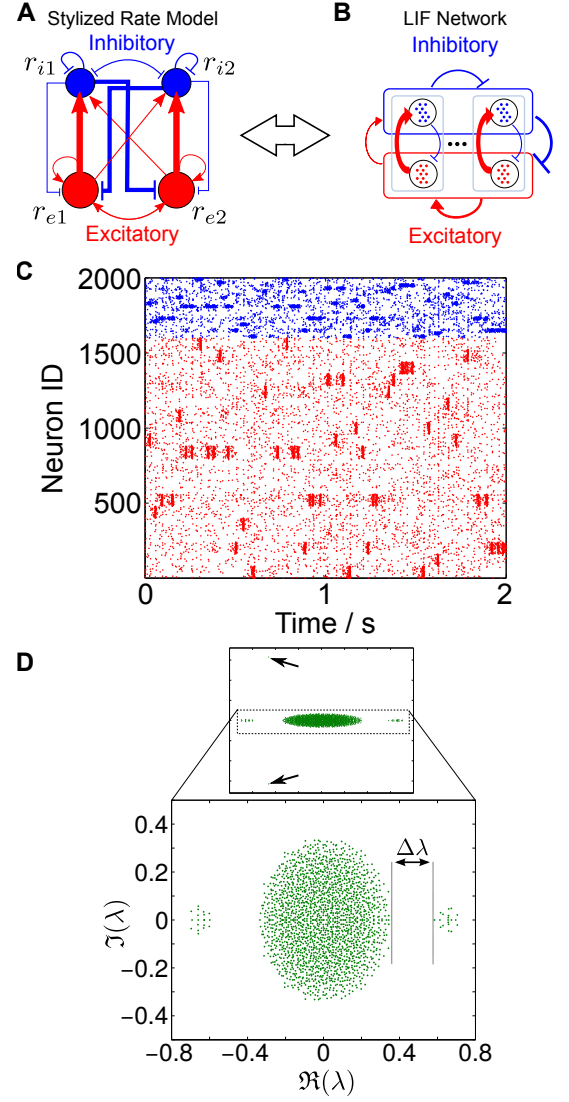


Figure 7. SSA dynamics can result from excitatory-to-inhibitory feedback loops. **A** Schematic of a stylized network model with two excitatory-to-inhibitory feedback loops corresponding to the model (8). **B** Illustration of the corresponding LIF network with such a co-clustered feedback mechanism between neuron types. **C** Raster plot of the dynamics of a co-clustered LIF network with $N = 2000$ neurons and $c = 20$ pairs with excitatory-to-inhibitory feedback. Note that the inhibitory neurons also exhibit SSA dynamics here. **D** Spectrum of the weight matrix exhibiting a spectral gap both on the left and right hand side of the bulk of the spectrum (see text for details).

lowing Schur form

$$Q = \begin{pmatrix} -w^+ & w_{ff} & & \\ & 0 & & \\ & & \sqrt{k}(s - \epsilon) & w_{ff,2} \\ & & -\sqrt{k}(s - \epsilon) & \end{pmatrix}, \quad (10)$$

$$w^+ = (k - 1)(s + \epsilon),$$

$$w_{ff} = (k + 1)(s + \epsilon), \quad w_{ff,2} = -(k - 1)(s - \epsilon),$$

where the eigenvalues of W are on the diagonal. When the leak term is considered, the largest eigenvalue of (1) is $(-1 + \sqrt{k}(s - \epsilon))$. Hence, to keep the system stable we need to constrain

$$0 < \sqrt{k}(s - \epsilon) < 1, \quad (11)$$

and the spectral gap is again controlled by $s - \epsilon$. In the associated orthonormal Schur basis, the first two modes of the firing rate dynamics

$$\mathbf{u}_1 = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \quad \mathbf{u}_2 = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ -1 \\ -1 \end{pmatrix} \quad (12)$$

are again global ‘sum’ and ‘difference’ modes, which interact via a balanced amplification mechanism [24]. However, there are also two localized modes \mathbf{u}_3 and \mathbf{u}_4 that can lead to slow structured activity:

$$\mathbf{u}_3 = \frac{1}{\sqrt{2k+2}} \begin{pmatrix} \sqrt{k} \\ -\sqrt{k} \\ 1 \\ -1 \end{pmatrix}, \quad \mathbf{u}_4 = \frac{1}{\sqrt{2k+2}} \begin{pmatrix} 1 \\ -1 \\ -\sqrt{k} \\ \sqrt{k} \end{pmatrix}. \quad (13)$$

Of these, \mathbf{u}_3 is associated with the largest eigenvalue and describes the slow(est) dynamics. This mode corresponds to a firing pattern of correlated activity within the pairs (r_{e1}, r_{i1}) and (r_{e2}, r_{i2}) , and anti-correlated activity across the pairs. As before, this analysis extends to networks with more groups and/or stochastic coupling between groups (for related arguments see supplementary material of Ref. [24]).

In Figure 7C-D, we show that the implementation of this wiring mechanism into a full LIF network displays SSA dynamics with a spectral gap, as predicted by our simple rate model, and the paired excitatory/inhibitory neurons act as a functional circuit displaying synchronous firing behavior. Importantly, note that, in contrast to the excitatory clustered scenario, there are no groups with dense reciprocal couplings in this network topology. Our numerics also show that varying the strength of the functional grouping ($W_{EI} = W_{IE}$) leads to the emergence of SSA dynamics, linked to the presence of the spectral gap $\Delta\lambda$ (Fig. 8A), and to the alignment of the leading Schur vectors with the ‘correct’ functional circuits, i.e., pairs of excitatory and inhibitory groups together (Figure 8B).

It is interesting to note that in this topology there is also a group of eigenvalues bounded away in the negative direction (see Figure 7D). These modes are the quickest decaying and correspond to ‘anti-correlated’ firing states, in which excitatory neurons act in synchrony with the inhibitory groups to which they are not functionally associated. Such fast decay reinforces the survival of synchrony within the functional groups in the network. Interestingly, these modes were already present in our linear rate model: \mathbf{u}_4 in (13) shows the same firing pattern with the fastest eigenvalue $\lambda_4 = -\sqrt{k}(s - \epsilon)$.




Figure 8. The presence of SSA dynamics in LIF networks with excitatory to inhibitory feedback loops. **A** The spike-rate variability, which measures the intensity of SSA dynamics, as a function of $W_{EI} = W_{IE}$ (left; dots: raw data from simulations, line: mean, shading: standard deviation) and $\Delta\lambda$ (right). See text and Fig. 4 for further details. **B** The first principal angle between the subspaces of the principal firing patterns and the dominant Schur vectors of the weight matrix W show high alignment. See text and Fig. 5 for further details.

D. SSA dynamics in LIF networks with alternative connectivity topologies

Although perhaps the most intuitive way of generating SSA dynamics follows from clustering the connectivity of the LIF network, our analysis above suggests that alternative types of structural organization support SSA dynamics in LIF networks, as long as they are characterized by a gap in their eigenvalue spectrum and a measure of block-localization of the Schur vectors. More generally, one could conjecture that any low rank perturbation of the weight matrix of a balanced network which leads to an eigenvalue gap might be a valid candidate for a mechanism to generate SSA in neuronal dynamics. A few such network architectures are worth highlighting, as they have been considered of particular relevance in a neuroscience context.

1. SSA dynamics in networks with small-world organization

The first noteworthy example is the broad class of small-world like networks [41], since small-world organizations have been observed in many structural and functional neurophysiological networks [42]. While many modular networks can have the small-world property [43], here we focus on the archetypal small-world structure à la Watts and Strogatz [41], which does not display a distinct modular organization but instead may be seen as

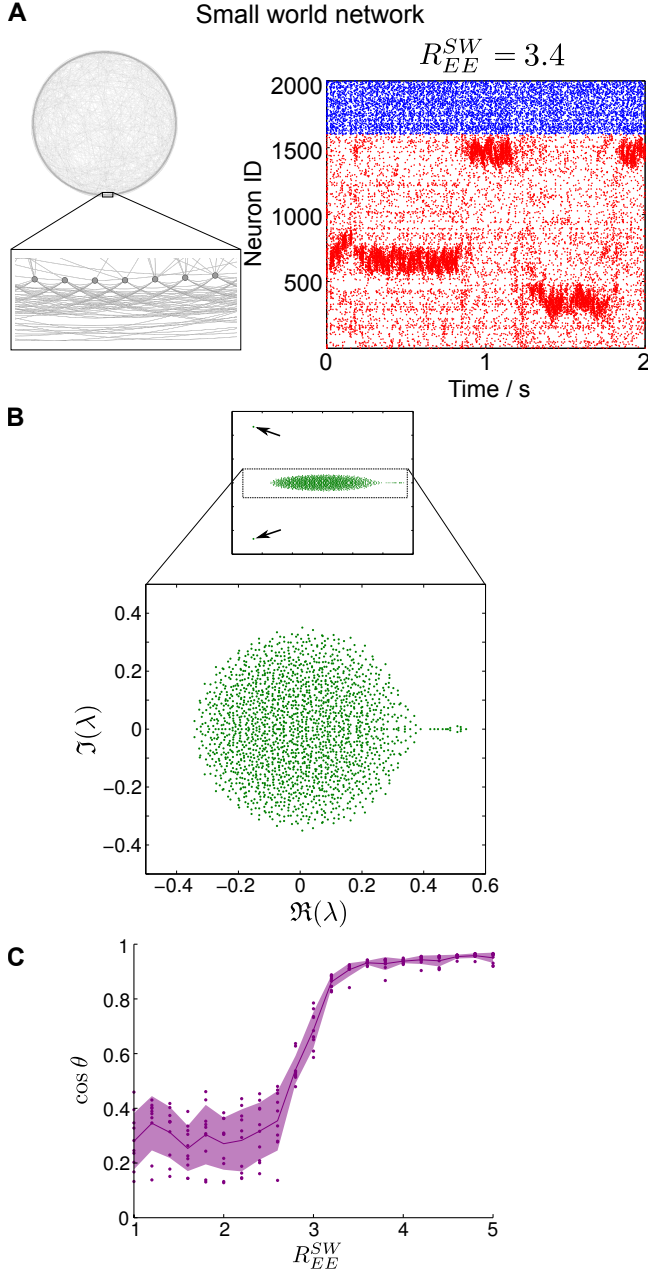


Figure 9. Small-world networks can exhibit SSA activity. **A** Small-world network and the resulting raster plot of its LIF network simulation. SSA activity is observed, yet with less distinct boundaries when compared to the clustered case. **B** Spectrum of the weight matrix of the small-world network showing a small eigenvalue gap. **C** The principal angle between the dominant subspace spanned by principal components of the firing rates and the dominant Schur vectors of the weight matrix shows that, as the modularity of the small-world rewiring becomes larger, there is alignment between the observed dynamics and the slow directions of the weight matrix (dots: raw data from simulations; line: mean; shading: standard deviation).

a mixture of a k -nearest neighbor ring lattice and a random graph. It is well known that such small-world net-

works possess distinctive spectral properties which have important implications for dynamical processes, such as synchronization [44]. As Figure 9A shows, small-world like LIF networks (see Materials and Methods for details of the construction) can indeed display SSA dynamics. In line with our arguments above, the weight matrix has a leading group of eigenvalues separated from the bulk, which dictate the slow dynamics, although in this case the gap is not as large and the slow eigenvalues are not as tightly bunched (Figure 9B). These separated eigenvalues correspond to a subspace spanned by eigenmodes with a slow spatial variation along the ‘backbone’ ring, and which lead to the localization and subsequent switching of the neural activity between subgroups of neurons. The principal angle between the firing patterns and the dominant Schur vectors (Figure 9C) indicates that the firing patterns are again dictated by the spectral properties of the underlying small-world topology. Note, however, that due to the lack of hard boundaries in the spectral groupings, both for the eigenvalue structure and the smoothly-varying dominant Schur vectors, the SSA dynamics (Figure 9A), although present, is not as distinctly localized as in the examples above. This behavior is effectively a result of the heavy overlap induced by the underlying lattice-like pristine world in the small-world construction [44].

2. Lack of SSA dynamics in scale-free networks

Another class of networks that has attracted tremendous interest is so-called scale-free networks [45]. These networks are characterized by a fat-tailed degree distribution, a fact that has been observed empirically for a variety of networks. As shown for undirected random graphs, the spectrum of scale-free networks may have a spectral gap, due to the presence of hubs with very large degrees [27]. However, the associated eigenvectors are usually localized around the hubs, and as such are not able to trigger a consistent pattern of localized activity across a group of nodes. In our simulations of scale-free LIF networks, we did not observe SSA dynamics: the dynamics was essentially concentrated around the hub with no transitions in time (*i.e.*, only a core of neurons around the hub had consistently elevated firing rates), hence no SSA. Indeed, the networks had only one large eigenvalue separated from the bulk (due to the high degree of the main hub) and its associated Schur vector was highly localized around the hub [27]. An example of this hub-centric dynamics can be found in Figure S2 of the Supplementary Information.

3. SSA dynamics with multiple time scales in LIF networks

It is also possible to enforce hierarchical arrangements of the functional wiring mechanisms discussed so far, thus allowing for a variety of combinatorial arrangements. De-

pending on the number of hierarchical layers, this enables the introduction of *multiple time scales* in the spatio-temporal segregated activity. As an illustrative example, Figure 10 shows the dynamics of a LIF network with a two-level hierarchy of clustered excitatory neurons. This network exhibits SSA with two ‘slow’ time scales: a very slow switching between the groups of the top level of the hierarchy, and the not-so-slow pulsation of activity between the subgroups within each of the large groups, thus reflecting the second level of the hierarchy. We remark that one could have employed a combination of different ‘functional circuits’ in the individual layers of the hierarchy leading to potentially different behaviors. As real neural networks display multiple layers of hierarchical organization [15–18, 46], understanding such schemes is of great importance for neural computation, and will be addressed in future work.

III. DISCUSSION

Answering the question of how the wiring of neural circuits governs neural dynamics is a key step in understanding neuronal computations. Here, we showed how knowledge of certain spectral features of the matrix of neuronal connectivity can help understand what dynamics a network can support.

In this paper we focused on LIF networks exhibiting slow-switching assembly (SSA) dynamics, where groups of neurons show a sustained and coherent increase in their firing rates, with slow switching between epochs of localized firing in different groups across the network. We found that the presence of the slow switching time scale is reflected in the spectral properties of the synaptic weight matrix: a gap separating the leading eigenvalues together with a block-localization of the associated Schur vectors on groups of neurons is a key indicator of the presence of SSA dynamics in the network. In line with this observation, multiple gaps in the eigenvalue spectrum are indicative of further time scales in the network dynamics (see Figure 10). Moreover, when the leading eigenvalue becomes larger than one, the SSA dynamics becomes localized on one of the assemblies (see Figure 6).

First, we revisited the case of balanced LIF networks with clustered excitatory neurons [22] and observed that only when there is an eigenvalue gap and Schur block-localization does the network display SSA dynamics. Further analytical understanding from a stylized firing-rate model allowed us to determine that the clustering strength drives the development of the eigenvalue gap responsible for the slow switching between localized firing modes. We remark that clustered excitatory connectivity leads to a Hebbian amplification regime [24]: the more stable the pattern formation, the slower the dynamics of any banded activity. Fast switching between well-defined, stable patterns is thus only achievable for strong input changes.

As suggested from our spectral characterization, and

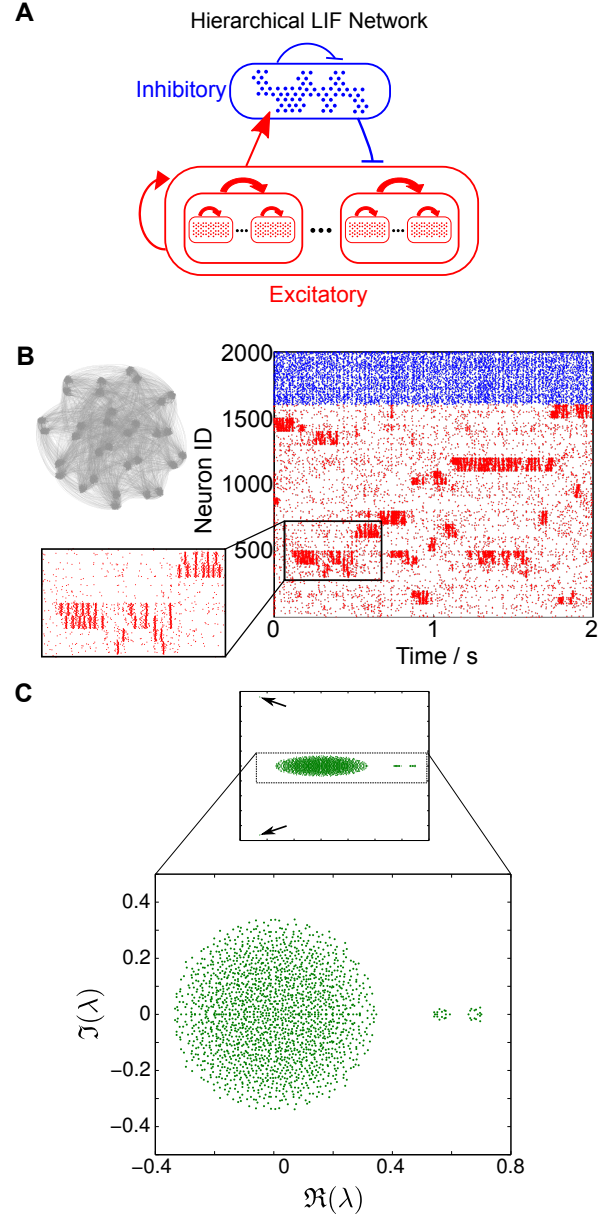


Figure 10. SSA activity in hierarchical LIF networks. **A** Schematic of a hierarchically clustered LIF network. Each excitatory group consists of two subgroups. **B** A hierarchical modular LIF network with $N = 2000$ neurons and the resulting raster plot of its simulation. SSA activity is observed at two time scales, corresponding to the two hierarchical levels embedded in the network structure: slow switching between the large groups and faster switching between the inner subgroups (see inset). **C** The spectrum of the weight matrix of this hierarchical network exhibits two eigenvalue gaps corresponding to the two slow time-scales the network can support.

confirmed by simulations, we showed that SSA dynamics can be achieved through the structured tuning of synaptic strengths, rather than through the clustered rewiring of the connections. While the equivalence between topological and weight organizations is clear for a linear sys-

tem, such a direct correspondence is not guaranteed a priori for a non-linear system. In fact, the clustering of weights appears to have a slightly larger influence compared to topological organization in our simulations of non-linear LIF systems. More importantly, this dynamical equivalence between clustered topology and clustered synaptic weights has potential ramifications for learning and synaptic plasticity; it indicates that the alteration of the weight structure can lead to the emergence of grouped activity without the need for structural rewiring of physical synaptic connections, thus suggesting a cost-effective adaptation to stochastically encode different firing patterns [47].

Within LIF networks with clustered excitatory connectivity [22], inhibitory neurons only play a balancing background role, whereas experimental studies have revealed a vast diversity of interneuron subtypes that play key roles in computation [48, 49]. We thus investigated mechanisms in which inhibitory neurons have an active, functional role in generating SSA dynamics. We demonstrated both analytically and via numerical simulations how such functional circuits can be constructed with the help of a feedback loop between excitatory and inhibitory neurons, so that inhibitory neurons themselves exhibit SSA dynamics and become an integral functional part of the dynamics. As fewer inhibitory neurons are present, adapting the weights according to this functional co-clustering scheme may also provide a cost effective alternative to generate SSA activity in neuronal networks.

The emergence of SSA dynamics can be achieved not only by clustering excitatory neurons but also by shaping the network structure in several ways. For instance, a small-world network topology of the excitatory neurons is also able to support SSA dynamics due to the spectral properties of small-worlds, which induce a separation of the slow eigenvalues and the localization and switching of firing activity associated with slowly spatially-varying dominant Schur vectors. In contrast, scale-free networks lack the block-localization of the Schur vectors on multiple groups which is necessary to consistently induce SSA dynamics in LIF networks. Further topologies will be explored in the future. In particular, inhibitory neurons that inhibit other inhibitory neurons (so-called disinhibition patterns [48, 50–54]) provide an interesting example. A different avenue might be provided by balanced amplification mechanisms [24], which could be potentially used for creating grouped activity, yet without the introduction of a slow time scale. Other interesting questions in this respect are: which wiring mechanisms provide the most economical, or evolutionarily fit, variant [47, 55] to induce a given dynamics; and how does the observed diversity of interneurons (and the multiple roles they can play) relate to this dynamical picture.

Our work emphasizes the importance of the spectral characterization of the weight matrix for the dynamics taking place on these topologies. Although spectral properties are also key in characterizing the dynamical response of networks operating at criticality [56], our ob-

servations here correspond to a different phenomenon. The emergence of a dominant assembly when $\lambda_{\max} \gtrsim 1$ leads to a saturation of the network and a decrease of its dynamical heterogeneity, in contrast to systems of coupled non-modular excitatory units at criticality, which maximize their dynamic range for $\lambda_{\max} = 1$ [56, 57].

Our work links up with experimental findings that cortical states can arise spontaneously and exhibit switching behavior. Recordings from anesthetized cat visual cortex have revealed activity states that dynamically switch, and these states match closely to recorded orientation maps [58]. Such cortical states likely arise due to similarly tuned neurons having higher intra-cortical connectivity [11, 12], as is also predicted through models of orientation maps [59–61]. These observations are similar to our theoretical outcomes which indicate that dynamic transitioning between different network states can be driven and sustained based only on the underlying network topology. Other experiments have identified states in CA3 network activity of the hippocampus [62] where active cell assemblies exist for tens of seconds before sharply transitioning to a new state. However, these were metastable states (hence unlikely to be reactivated) and included a core population of cells consistently active in all states. While such complex dynamics are beyond the simple models presented here, it may be feasible to model such metastable dynamics through more elaborate network topologies, which may include synaptic plasticity [63]. Finally, let us remark that our reduced model descriptions effectively focused on networks implementing a *rate-based* coding mechanism and do not include spike-timing of cell assemblies, which has been identified to be an important component in neuronal computation, e.g. in rodents [64, 65], monkeys [66], songbirds [67], and grasshoppers [68]. Whether our results can be translated to a time-coding regime would be an interesting question for future work.

As connectomics continues to advance the mapping of connections and synaptic strengths in wide areas of the brain, experimentally obtained weight matrices can be analyzed spectrally as described above to determine if SSA dynamics can be supported by the networks under study. While we focused here on the simplest mechanisms underlying SSA activity, future work could also consider the construction of biophysically realistic models that take into account the growing literature on the distribution of synaptic strengths, synaptic contacts, firing rates, and other relevant cortical parameters that tend to show a lognormal distribution [7, 69–71].

Although knowledge of the network structure (connectomics) is not sufficient to predict the dynamics a network circuit will display (for instance, the network dynamics can be dominated completely by a strong input), it can still give valuable insights on the firing patterns the network can support. Conversely, the observation of neuronal dynamics alone may not be sufficient to understand neural computation in detail, since different network topologies can yield similar dynamics. Hence, our

work hints at how connectomics and neuronal dynamics data can provide complementary and intertwined routes for systems neuroscientists to study the computational principles implemented by the brain.

Finally, it is important to remark that in order to get a fuller picture of the relation between structure and dynamical network properties many other factors not considered here will be of interest, including the distribution of inputs, the precise location of synapses on the post-synaptic neuron, and the plasticity rules that govern the evolution of the network over time [72]. Linking the insight gained from our work to experimental observations or to highly detailed computational models [3, 73] would thus be a fruitful next step in order to bridge the gap between neuronal structure, dynamics, and ultimately function.

IV. MATERIAL AND METHODS

A. Leaky integrate-and-fire networks

We simulated leaky-integrate-and-fire (LIF) networks where the non-dimensionalized membrane potential of each neuron ($V_i(t)$, $i = 1, \dots, N$) was modeled by:

$$\frac{dV_i(t)}{dt} = \frac{1}{\tau_m}(\mu_i - V_i(t)) + \sum_j W_{ij}g_j^{E/I}(t), \quad (14)$$

with a firing threshold of 1 and a reset potential of 0. The constant input terms μ_i were chosen uniformly in the interval $[1.1, 1.2]$ for excitatory neurons, and in the interval $[1, 1.05]$ for inhibitory neurons. The membrane time constants for excitatory and inhibitory neurons were set to $\tau_m = 15$ ms and $\tau_m = 10$ ms, respectively. The refractory period was fixed at 5 ms for both excitatory and inhibitory neurons. Note that although the constant input term is supra-threshold, balanced inputs guaranteed that the average membrane potential is sub-threshold [22].

The network dynamics is captured by the sum in (14), which describes the input to neuron i from all other neurons in the network. The topology of the network is encoded by the weight matrix W , i.e., W_{ij} denotes the weight of the connection from neuron j to neuron i , where W_{ij} is zero if there is no connection. Synaptic inputs are modeled by $g_j^{E/I}(t)$, which is increased step-wise instantaneously after a presynaptic spike of neuron j ($g_j^{E/I} \rightarrow g_j^{E/I} + 1$) and then decays exponentially according to:

$$\tau_{E/I} \frac{dg_j^{E/I}}{dt} = -g_j^{E/I}(t), \quad (15)$$

with time constants $\tau_E = 3$ ms for an excitatory interaction, and $\tau_I = 2$ ms if the presynaptic neuron is inhibitory.

Equation (14) can be rewritten in matrix notation as:

$$\frac{d\mathbf{V}(t)}{dt} = T^{-1}[\boldsymbol{\mu} - \mathbf{V}] + W\mathbf{g}(t), \quad (16)$$

where $T = \text{diag}(\tau_i)$ and $\mathbf{V} = [V_1, \dots, V_N]^T$ and the vectors \mathbf{V} , $\boldsymbol{\mu}$, and \mathbf{g} are $N \times 1$ vectors. The spectral analyses in the main text (eigenvalues and Schur vectors) refer to the weight matrix W . Different network topologies correspond to different weight matrices W , as explained below.

In all simulations, the ratio of excitatory to inhibitory neurons was fixed to be 4 : 1. Unless otherwise stated, the networks comprise $N = 2000$ units (1600 excitatory, 400 inhibitory) and were simulated over 20 seconds to calculate the statistics reported. The time step for all simulations was 0.1 ms. All simulations were run in MATLAB (version 2011b or later).

B. Network Topologies and Weight matrices

Different network topologies were simulated using different W matrices but maintaining a general balanced network. If not indicated differently, all parameters below correspond to the case $N = 2000$.

1. Networks with unclustered balanced connections

Unless otherwise stated, unclustered connections between neurons were drawn uniformly at random according to the probabilities $p^{EE} = 0.2$ and $p^{EI} = p^{IE} = p^{II} = 0.5$, where the first superscript denotes the destination and the second superscript denotes the origin of the synaptic connection, and E, I stand for an excitatory or inhibitory neuron, respectively. Synaptic weights between inhibitory neurons had always the weight $w^{II} = -0.0297$. The other weight parameters were set to $w^{EI} = -0.0297$, $w^{IE} = 0.0074$ and $w^{EE} = 0.0156$, except where indicated differently in the scenarios described below.

2. Networks with clustered excitatory-to-excitatory connections

First, LIF networks with clustered excitatory neurons were constructed according to the protocol in Ref. [22]. The connection probabilities between excitatory neurons were changed so that neurons had more connections within its group than to neurons outside the group. The average number of connections was kept the same as for an unclustered balanced network ($p^{EE} = 0.2$) while varying the ratio $R_{EE} = p_{in}/p_{out}$, where p_{in} and p_{out} are the probabilities of connectivity within a group and between groups, respectively. $R_{EE} = 1$ corresponds to the unclustered balanced network.

Second, we generated LIF networks where the connectivity between all excitatory neurons was uniform ($p^{EE} = 0.2$, $R_{EE} = 1$) but the weights were varied by changing the ratio $W_{EE} = w_{in}^{EE}/w_{out}^{EE}$, where w_{in}^{EE} refers to synaptic weights within the same group and w_{out}^{EE} refers to synaptic weights to neurons in other groups.

The average synaptic weight between excitatory neurons was kept at $w^{EE} = 0.0156$, as for the unclustered case. Strictly speaking, this scheme does not correspond to a clustering in a topological sense, but rather to a clustering of synaptic weights between excitatory neurons. However, in terms of the dynamics, both E-E clustered network variants (adjusting probabilities or weights) lead to effectively the same results, as discussed in the text. Unless otherwise stated, the excitatory neurons were divided into 20 groups of 80 neurons each.

As the weights in balanced networks should scale with $1/\sqrt{N}$ [22, 23], when varying the network size (see Fig. 2A) the parameter settings were kept the same but all weights were scaled accordingly, e.g., for $N = 1000$ the weights were multiplied by $\sqrt{2}$ compared to those of the network with $N = 2000$.

3. Networks with excitatory-to-inhibitory feedback loop

To construct LIF networks with co-clustered excitatory and inhibitory neurons, we kept uniform excitatory-to-excitatory and inhibitory-to-inhibitory couplings while varying the ratios $W_{IE} = w_{in}^{IE}/w_{out}^{IE}$, $W_{EI} = (w_{in}^{EI}/w_{out}^{EI})^{-1}$, or the analogous ratios of connections probabilities R_{IE}, R_{EI} . The subscript ‘in’ indicates connections within the excitatory/inhibitory pair. Note that W_{EI}, R_{EI} are defined by the inverse in/out ratio, since they describe an inhibitory effect. To modulate the co-clustered excitatory-to-inhibitory network dynamics, we vary the ratios W_{IE}, W_{EI} , while keeping the average weights constant. For instance, in the example shown in Figure 8, we kept $R_{EI} = R_{IE} = 2$ while simultaneously varying the ratios $W_{IE} = W_{EI}$ between 1 and 5. For all networks, we divided the network into 20 groups with 80 excitatory and 20 inhibitory neurons each.

4. Networks with small-world connectivity

Networks with small-world connectivity between excitatory neurons were constructed as follows. Excitatory neurons ($N = 1600$) were connected within a 40-nearest neighbor ring with probability p_{in}^{EE} , i.e., neurons i was connected to neuron j with probability p_{in}^{EE} if $|i - j| \leq 40$ subject to periodic boundary conditions. The connection probability outside this neighborhood was set to p_{out}^{EE} . As in previous models, the average connectivity was kept constant at $p^{EE} = 0.2$ while varying the ratio $R_{EE}^{SW} = p_{in}^{EE}/p_{out}^{EE}$. Hence, by increasing R_{EE}^{SW} we increase the amount of ‘backbone’ connectivity through the (stochastic) 40 nearest-neighbor cycle, while decreasing the number of connections to neurons elsewhere in the network. This construction may thus also be interpreted as an overlapping clustering and is effectively a variant of the small-world scheme introduced by Watts and Strogatz [41].

5. Networks with scale-free connectivity

LIF networks with a scale-free topology for the excitatory connections were created by a preferential attachment process [45] using algorithm 5 in [74] with parameter $d = 64$. The weight matrix was created iteratively by adding one neuron at a time, such that the probability that it is connected to an existing node is proportional to the degree of that node (at the current state), i.e., the earlier a node is added, the larger its final degree will be. The resulting degree of the network tends to a power-law distribution [45]. The value $d = 64$ was chosen to avoid saturation of the dynamics of the network, and this leads to an excitatory-excitatory matrix with 35% of the edges of our other examples.

6. Network with hierarchical excitatory connections

We simulated LIF networks where excitatory neurons belonged to a hierarchy of groups (Figure 10) by dividing the excitatory neurons into nested clusters. We implemented two layers in the hierarchy: the excitatory neurons were divided into 16 groups, and each of these groups was then subdivided into 2 more groups to give a total of 32 subgroups. This was achieved by varying the ratios $R_{EE}^{top} = p_{group}^{EE}/p_{out}^{EE}$, $R_{EE}^{sub} = p_{subgroup}^{EE}/p_{group}^{EE}$ while keeping the average connectivity between excitatory neurons constant. The example in Figure 10 corresponds to $R_{EE}^{top} = 1.45$ and $R_{EE}^{sub} = 3.7$. The weight within the subgroups was set at $w_{sub}^{EE} = 0.0163$. All other connections remained unchanged.

C. Quantifying SSA dynamics from spike-train LIF simulations

To evaluate the extent to which a network displays slow-switching assembly dynamics, we used the following two spike-rate variability measures.

Given a partition of the neurons into c groups, we compute the average spiking frequency $f_i(t)$ of the neurons in each cluster $i \in \{1, \dots, c\}$ over non-overlapping windows of 100ms. For each time window, we obtain the firing rate vector $\mathbf{f}(t) = [f_1(t), \dots, f_c(t)]^T$ of which we compute the standard deviation $\sigma(t)$. The standard deviations are then averaged over the duration of the simulation:

$$S = \frac{1}{T} \sum_{t=1}^T \sigma(t), \quad (17)$$

where T is the total number of time windows in the simulation. We then obtain a boot-strapped expectation $\langle S^{\text{shuff}} \rangle$ computed by reshuffling neurons at random into groups of the same sizes as those in the partition. The spike rate variability score is then defined as:

$$\hat{S} = S - \langle S^{\text{shuff}} \rangle, \quad (18)$$

where the average $\langle S^{\text{shuff}} \rangle$ is obtained over 10 random reshufflings of the neurons. If the network fires uniformly (with no localized patterns), \hat{S} is low; whereas \hat{S} increases if the network displays heterogeneous activity aligned with the partition under investigation.

As explained in the text, we introduced a second spike-train variability measure to quantify the variations of the group firing-rates across time. This complementary measure allows us to discern scenarios in which there is a group of neurons dominating the firing (thus leading to a large variation across groups), but no switching between groups. To quantify these effects we use an analogous measure to \hat{S} above.

Given a partition of the neurons into c groups, we compute the average spiking frequency $f_i(t)$ of the neurons in each cluster $i \in \{1, \dots, c\}$ over non-overlapping windows of 100ms. For each group i , we compute the standard deviation $\sigma_T(i)$ of the vector of coarse-grained firing rates across time $\mathbf{f}_i = [f_i(t_1), \dots, f_i(t_N)]$, where t_k stands for the k th time bin and $i = 1, \dots, c$. We then average over groups to obtain:

$$S_T = \frac{1}{c} \sum_{i=1}^c \sigma_T(i). \quad (19)$$

As above, we obtain a boot-strapped expectation $\langle S_T^{\text{shuff}} \rangle$ by reshuffling neurons at random into groups of the same sizes as those in the partition. The spike rate variability score over time is then defined as:

$$\hat{S}_T = S_T - \langle S_T^{\text{shuff}} \rangle, \quad (20)$$

where the average $\langle S_T^{\text{shuff}} \rangle$ is obtained over 10 random reshufflings of the neurons. Again, if the network fires homogeneously (with no localized patterns in time), \hat{S}_T is low; whereas \hat{S}_T increases if the network displays heterogeneous firing rates over time.

In Figure S1, we show the behavior of \hat{S}_T for the examples discussed in the main manuscript. As expected,

for all cases where SSA dynamics is present, both measures \hat{S} and \hat{S}_T behave consistently. On the other hand, \hat{S}_T detects the end of the SSA region when, through increased clustering, the dynamics gets localized on one cell assembly, as shown in Figure 6.

D. Measuring alignment of LIF dynamics with the Schur vectors of the weight matrix: the principal angle

First, we find the dominant firing patterns in the network dynamics. We perform simulations of the LIF network and obtain the firing rates of every neuron in 250 ms bins to generate an $N \times T$ matrix, where N is the number of neurons and T is the number of bins. On this matrix, we perform a *principal component analysis* (PCA) and select the first $c - 1$ principal components $\{\mathbf{p}_i\}_{i=1}^{c-1}$. This set of N -dimensional vectors captures most of the variability observed in the simulated dynamics.

We then assess how aligned the $c - 1$ principal components are with the dominant Schur vectors of the weight matrix of the network. More precisely, we compute the Schur decomposition of the weight matrix W ; keep the $c - 1$ dominant Schur vectors $\{\mathbf{u}_i\}_{i=1}^{c-1}$ associated with the eigenvalues with largest real part; and then compute the (first) principal or canonical angle θ between the subspaces spanned by the two sets of vectors $\mathcal{P} = \text{span}\{\mathbf{p}_i\}$ and $\mathcal{U} = \text{span}\{\mathbf{u}_i\}$ [38, 39]:

$$\cos(\theta) = \max \left\{ \frac{\mathbf{u}^T \mathbf{p}}{\|\mathbf{u}\| \|\mathbf{p}\|} \mid \mathbf{u} \in \mathcal{U} \ \mathbf{p} \in \mathcal{P} \right\}, \quad (21)$$

The first principal angle measures how ‘close’ the observed firing patterns are to the dominant modes computed solely from the weight matrix. If $\cos(\theta) \approx 1$, there is a large alignment between the span of both subspaces.

-
- [1] G. Buzsaki, *Nature Neuroscience* **7**, 446 (2004).
 - [2] J. Du, T. J. Blanche, R. R. Harrison, H. A. Lester, and S. C. Masmanidis, *PLoS ONE* **6** (2011), 10.1371/journal.pone.0026204.
 - [3] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller, *Nature Methods* **10**, 413+ (2013).
 - [4] G. Buzsaki, *Neuron* **68**, 362 (2010).
 - [5] D. O. Hebb, *The Organization of Behavior* (John Wiley & Sons, 1949).
 - [6] K. Harris, *Nature Reviews Neuroscience* **6**, 399 (2005).
 - [7] S. Song, P. Sjöström, M. Reigl, S. Nelson, and D. Chklovskii, *Plos Biology* **3**, 507 (2005).
 - [8] R. Perin, T. K. Berger, and H. Markram, *Proceedings of the National Academy of Sciences of the United States of America* **108**, 5419 (2011).
 - [9] Y. Yoshimura, J. Dantzer, and E. Callaway, *Nature* **433**, 868 (2005).
 - [10] T. Otsuka and Y. Kawaguchi, *Journal of Neuroscience* **31**, 3862 (2011).
 - [11] H. Ko, S. B. Hofer, B. Pichler, K. A. Buchanan, P. J. Sjoestrom, and T. D. Mrsic-Flogel, *Nature* **473**, 87 (2011).
 - [12] K. D. Harris and T. D. Mrsic-Flogel, *Nature* **503**, 51 (2013).
 - [13] S. Lefort, C. Tómm, J. C. F. Sarria, and C. C. H. Petersen, *Neuron* **61**, 301 (2009).
 - [14] L. Yassin, B. L. Benedetti, J.-S. Jouhanneau, J. A. Wen, J. F. A. Poulet, and A. L. Barth, *Neuron* **68**, 1043 (2010).
 - [15] M. Shimono and J. M. Beggs, *Cerebral Cortex* (2014), 10.1093/cercor/bhu252.
 - [16] M. J. McGinley and G. L. Westbrook, *Proceedings of the National Academy of Sciences of the United States*

- of America **110**, 16193 (2013).
- [17] I. Savic, B. Gulyas, M. Larsson, and P. Roland, *Neuron* **26**, 735 (2000).
 - [18] D. J. Felleman and D. C. Van Essen, *Cerebral Cortex* **1**, 1 (1991).
 - [19] M. Ito, N. Masuda, K. Shinomiya, K. Endo, and K. Ito, *Current Biology* **23**, 644 (2013).
 - [20] C. J. Pavlides, Y. Greenstein, M. Grudman, and W. J. Brain Research **439**, 383 (1988).
 - [21] J. Hyman, B. Wyble, V. Goyal, C. Rossi, and M. Hasselmo, *Journal of Neuroscience* **23**, 11725 (2003).
 - [22] A. Litwin-Kumar and B. Doiron, *Nature Neuroscience* **15**, 1498 (2012).
 - [23] C. van Vreeswijk and H. Sompolinsky, *Neural Computation* **10**, 1321 (1998).
 - [24] B. K. Murphy and K. D. Miller, *Neuron* **61**, 635 (2009).
 - [25] U. von Luxburg, *Statistics and Computing* **17**, 395 (2007).
 - [26] X. Zhang, R. Rao Nadakuditi, and M. E. J. Newman, "Spectra of random graphs with community structure and arbitrary degrees," (2013), arXiv:1310.0046.
 - [27] R. R. Nadakuditi and M. E. J. Newman, *Physical Review E* **87**, 012803 (2013).
 - [28] H. A. Simon and A. Ando, *Econometrica* **29**, 111 (1961).
 - [29] R. F. Galán, *PLoS ONE* **3**, e2148 (2008).
 - [30] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, *Proceedings of the National Academy of Sciences* **107**, 12755 (2010).
 - [31] M. T. Schaub, J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, *PLoS ONE* **7**, e32210 (2012).
 - [32] J.-C. Delvenne, M. T. Schaub, S. N. Yaliraki, and M. Barahona, in *Dynamics On and Of Complex Networks, Volume 2*, edited by A. Mukherjee, M. Choudhury, F. Peruani, N. Ganguly, and B. Mitra (Springer New York, 2013) pp. 221–242.
 - [33] Y. N. Billeh, M. T. Schaub, C. A. Anastassiou, M. Barahona, and C. Koch, *Journal of Neuroscience Methods* **236**, 92 (2014).
 - [34] K. Rajan and L. F. Abbott, *Physical Review Letters* **97**, 188104 (2006).
 - [35] H. J. Sommers, A. Crisanti, H. Sompolinsky, and Y. Stein, *Physical Review Letters* **60**, 1895 (1988).
 - [36] M. S. Goldman, *Neuron* **61**, 621 (2009).
 - [37] L. N. Trefethen and M. Embree, *Spectra and pseudospectra: the behavior of nonnormal matrices and operators* (Princeton University Press, 2005).
 - [38] H. Golub and F. C. Van Loan, *Matrix Computations*, 3rd ed. (John Hopkins University Press, Baltimore and London, 1996).
 - [39] G. W. Stewart, *Matrix Algorithms Volume 2: Eigensystems*, Vol. 2 (Siam, 2001).
 - [40] U. Rutishauser, R. J. Douglas, and J.-J. Slotine, *Neural Computation* **23**, 735 (2011).
 - [41] D. J. Watts and S. H. Strogatz, *Nature* **393**, 440 (1998).
 - [42] E. Bullmore and O. Sporns, *Nature Reviews Neuroscience* **10**, 186 (2009).
 - [43] D. Meunier, R. Lambiotte, and E. T. Bullmore, *Frontiers in Neuroscience* **4** (2010), 10.3389/fnins.2010.00200.
 - [44] M. Barahona and L. M. Pecora, *Physical Review Letters* **89**, 054101 (2002).
 - [45] A.-L. Barabási and R. Albert, *Science* **286**, 509 (1999).
 - [46] J. Ambrosingerson, R. Granger, and G. Lynch, *Science* **247**, 1344 (1990).
 - [47] S. Laughlin and T. Sejnowski, *Science* **301**, 1870 (2003).
 - [48] A. Kepecs and G. Fishell, *Nature* **505**, 318–326 (2014).
 - [49] L. Roux and G. Buzsáki, *Neuropharmacology*, (2014).
 - [50] G. Silberberg and H. Markram, *Neuron* **53**, 735 (2007).
 - [51] H. Xu, H.-Y. Jeong, R. Tremblay, and B. Rudy, *Neuron* **77**, 155 (2013).
 - [52] S. Lee, I. Kruglikov, Z. J. Huang, G. Fishell, and B. Rudy, *Nature Neuroscience* **16**, 1662 (2013).
 - [53] H.-J. Pi, B. Hangya, D. Kvitsiani, J. I. Sanders, Z. J. Huang, and A. Kepecs, *Nature* **503**, 521 (2013).
 - [54] C. K. Pfeffer, M. Xue, M. He, Z. J. Huang, and M. Scanziani, *Nature Neuroscience* **16**, 1068 (2013).
 - [55] J. E. Niven and S. B. Laughlin, *Journal of Experimental Biology* **211**, 1792 (2008).
 - [56] D. B. Larremore, W. L. Shew, and J. G. Restrepo, *Physical Review Letters* **106** (2011), 10.1103/PhysRevLett.106.058101.
 - [57] J. Beggs and D. Plenz, *Journal of Neuroscience* **23**, 11167 (2003).
 - [58] T. Kenet, D. Bibitchkov, M. Tsodyks, A. Grinvald, and A. Arieli, *Nature* **425**, 954 (2003).
 - [59] R. Ben-Yishai, R. Bar-Or, and H. Sompolinsky, *Proceedings of the National Academy of Sciences of the United States of America* **92**, 3844 (1995).
 - [60] D. Somers, S. Nelson, and M. Sur, *Journal of Neuroscience* **15**, 5448 (1995).
 - [61] U. Ernst, K. Pawelzik, C. Sahar-Pikielny, and M. Tsodyks, *Nature Neuroscience* **4**, 431 (2001).
 - [62] T. Sasaki, N. Matsuki, and Y. Ikegaya, *Journal of Neuroscience* **27**, 517 (2007).
 - [63] A. Litwin-Kumar and B. Doiron, *Nature Communications* **5** (2014).
 - [64] K. Harris, J. Csicsvari, H. Hirase, G. Dragoi, and G. Buzsáki, *Nature* **424**, 552 (2003).
 - [65] D. Foster and M. Wilson, *Nature* **440**, 680 (2006).
 - [66] A. Riehle, S. Grun, M. Diesmann, and A. Aertsen, *Science* **278**, 1950 (1997).
 - [67] R. Hahnloser, A. Kozhevnikov, and M. Fee, *Nature* **419**, 65 (2002).
 - [68] A. Rokem, S. Watzl, T. Gollisch, M. Stemmler, A. Herz, and I. Samengo, *Journal of Neurophysiology* **95**, 2541 (2006).
 - [69] T. Hromádka, M. R. DeWeese, and A. M. Zador, *PLoS Biology* **6**, 124 (2008).
 - [70] G. Buzsáki and K. Mizuseki, *Nature Reviews Neuroscience* **15**, 264 (2014).
 - [71] S. W. Oh, J. A. Harris, L. Ng, B. Winslow, N. Cain, S. Mihalas, Q. Wang, C. Lau, L. Kuan, A. M. Henry, M. T. Mortrud, B. Ouellette, T. N. Nguyen, S. A. Sorensen, C. R. Slaughterbeck, W. Wakeman, Y. Li, D. Feng, A. Ho, E. Nicholas, K. E. Hirokawa, P. Bohn, K. M. Joines, H. Peng, M. J. Hawrylycz, J. W. Phillips, J. G. Hohmann, P. Wohnoutka, C. Koch, A. Bernard, C. Dang, A. R. Jones, H. Zeng, and C. R. Gerfen, *Nature* **508**, 207 (2014).
 - [72] C. Clopath, L. Buesing, E. Vasilaki, and W. Gerstner, *Nature Neuroscience* **13**, 344 (2010).
 - [73] M. W. Reimann, C. A. Anastassiou, R. Perin, S. L. Hill, H. Markram, and C. Koch, *Neuron* **79**, 375 (2013).
 - [74] V. Batagelj and U. Brandes, *Phys. Rev. E* **71**, 036113 (2005).

**Appendix A: Supplementary Information –
Emergence of slow-switching assemblies in
structured neuronal networks**




Figure S1. Spike rate variability across time, as a function of topological changes of the network. **A** Clustering Strength R_{EE} vs spike rate variability across time \hat{S}_T for varying network sizes with a fixed assembly size. Compare to Fig. 2A of the main text. **B** Clustering Strength R_{EE} vs spike rate variability across time \hat{S}_T for a fixed network size with varying assembly sizes. Compare to Fig. 2B of the main text. **C** Weight Clustering W_{EE} vs spike rate variability across time \hat{S}_T for a LIF network with 2000 neurons (20 groups; topological clustering $R_{EE} = 1$). Compare to Fig. 4 of the main text. **D** Spike rate variability across time \hat{S}_T as a function of $W_{EI} = W_{IE}$ for a network with excitatory to inhibitory feedback loops. Compare to Fig. 8 of the main text. In all plots dots denote raw data from simulations; line: mean, shading: standard deviation.




Figure S2. Network dynamics for a scale-free network. **A** Raster plot for an example scale-free network. Excitatory neurons are ordered according to their (expected) degree: high (bottom) to low (top). Note that the spiking activity is largely concentrated around the hub (which due to assortativity is connected to other nodes of high degree), and there are no distinguishable groups or switching behavior. **B** Spectrum of the simulated scale-free network. There is only one eigenvalue separated from the main bulk of the spectrum due to the presence of the large hub in the network.